

Adaptive sequential Monte Carlo by means of mixture of experts

J. Cornebise*, E. Moulines†, J. Olsson‡

August 12th, 2011

Abstract

Selecting appropriately the proposal kernel of particle filters is an issue of significant importance, since a bad choice may lead to deterioration of the particle sample and, consequently, waste of computational power. In this paper we introduce a novel algorithm approximating adaptively the so-called optimal proposal kernel by a mixture of integrated curved exponential distributions with logistic weights. This family of distributions is broad enough to be used in the presence of multi-modality or strongly skewed distributions. This "mixture of experts" is fitted, via Monte Carlo EM or online-EM methods, to the optimal kernel through minimization of the Kullback-Leibler divergence between the auxiliary target and instrumental distributions of the particle filter. The algorithm requires only one optimization problem to be solved for the whole sample, as opposed to existing methods solving one problem per particle. In addition, we illustrate in a simulation study how the method can be successfully applied to optimal filtering in nonlinear state-space models.

Keywords: Optimal and backward kernel, Adaptive algorithms, Kullback-Leibler divergence, Coefficient of variation, Expectation Maximization, particle filter

1 Introduction

During the last decade, sequential Monte carlo (SMC) methods have developed continuously from being a tool for solving rather specific problems, such as the optimal filtering problem in general state-space models or simulation of rare events

*Department of Statistical Science, University College London, julien@stats.ucl.ac.uk. J.C. is supported by grant BBSRC grant BB/G006997/1.

†Département Traitement du Signal et de l'Image, Telecom ParisTech

‡Centre for Mathematical Sciences, Lund University

(these topics are abundantly covered in [Liu, 2001](#); [Doucet et al., 2001](#); [Ristic et al., 2004](#); [Cappé et al., 2005](#), and the references therein), to a tool for sampling sequentially more or less *arbitrary* sequences of distributions ([Del Moral et al., 2006](#)). The high flexibility of these methods makes it possible to use them for executing efficiently key steps of composite algorithms tackling very complex problems, e.g., as in particle Markov chain Monte Carlo (MCMC) methods for joint inference on static and dynamic variables (see [Andrieu et al., 2010](#)).

However, while lots of research has been devoted to adaptation of MCMC methods (see especially [Haario et al., 2001](#); [Roberts and Rosenthal, 2007](#); [Andrieu and Moulines, 2006](#); [Roberts and Rosenthal, 2009](#)), adaptive methods for SMC is still in its infancy. This article, which sets out from theoretical framework for adaptive SMC presented in [Cornebise et al. \(2008\)](#), is part of the effort to fill this gap, which is mandatory for achieving maximum efficiency and user-friendliness.

In the SMC framework, the objective consists in approximating a *sequence* of target densities evolving sequentially by updating recursively a set of particles and associated importance weights. In order to describe precisely such an update, let ν be a target probability density function over a space $\Xi = \mathbb{R}^d$ and suppose that we have at hand a weighted sample $\{(\mathbf{X}_i, \omega_i)\}_{i=1}^N$ such that $\Omega^{-1} \sum_{j=1}^N \omega_j f(\mathbf{X}_j)$, with $\Omega = \sum_{j=1}^N \omega_j$, approximates $\int f(\mathbf{x}) \nu(\mathbf{x}) d\mathbf{x}$ for any ν -integrable function f . As a notational convention, we use capital letters for random variables, normal case for function arguments, and bold case for vectors.

We now wish to transform (by moving the particles and adjusting accordingly the importance weights) the sample $\{(\mathbf{X}_i, \omega_i)\}_{i=1}^N$ into a new weighted particle sample approximating some probability density μ on another space $\tilde{\Xi}$ given by

$$\mu(\tilde{\mathbf{x}}) := \frac{\int \nu(\mathbf{x}) l(\mathbf{x}, \tilde{\mathbf{x}}) d\mathbf{x}}{\iint \nu(\mathbf{x}) l(\mathbf{x}, \tilde{\mathbf{x}}) d\mathbf{x} d\tilde{\mathbf{x}}}, \quad (1.1)$$

where l is an un-normalized transition density. Here $\tilde{\Xi}$ can either be equal to Ξ or be a space of different dimension. Having access to $\{(\mathbf{X}_i, \omega_i)\}_{i=1}^N$, an approximation of μ is naturally obtained by plugging the weighted empirical measure associated with this weighted sample into (1.1), yielding the mixture density

$$\mu(\tilde{\mathbf{x}}) \approx \frac{\sum_{i=1}^N \omega_i l(\mathbf{X}_i, \tilde{\mathbf{x}})}{\sum_{j=1}^N \omega_j \int l(\mathbf{X}_j, \tilde{\mathbf{x}}) d\tilde{\mathbf{x}}} = \sum_{i=1}^N \left(\frac{\omega_i \int l(\mathbf{X}_i, \tilde{\mathbf{x}}) d\tilde{\mathbf{x}}}{\sum_{j=1}^N \omega_j \int l(\mathbf{X}_j, \tilde{\mathbf{x}}) d\tilde{\mathbf{x}}} \right) \left(\frac{l(\mathbf{X}_i, \tilde{\mathbf{x}})}{\int l(\mathbf{X}_i, \tilde{\mathbf{x}}) d\tilde{\mathbf{x}}} \right), \quad (1.2)$$

and ideally an updated particle sample would be obtained by drawing new particles $\{\tilde{\mathbf{X}}_i\}_{i=1}^N$ independently from (1.2). There are however two problems with this approach: firstly, in general, the integral $\int l(\cdot, \tilde{\mathbf{x}}) d\tilde{\mathbf{x}}$ lacks closed-form expression, ruling out direct computation of the mixture weights; secondly, even if the integral was known on closed-form, the resulting algorithm would have an $O(N^2)$ computational complexity due to the normalization of the mixture weights. To cope with these issues, we will, as proposed by [Pitt and Shephard \(1999\)](#), aim at

sampling instead the *auxiliary target distribution*

$$\mu^{\text{aux}}(i, \tilde{\mathbf{x}}) = \frac{\omega_i l(\mathbf{X}_i, \tilde{\mathbf{x}})}{\sum_{j=1}^N \omega_j \int l(\mathbf{X}_j, \tilde{\mathbf{x}}) d\tilde{\mathbf{x}}} = \frac{\omega_i a^*(\mathbf{X}_i)}{\sum_{j=1}^N \omega_j a^*(\mathbf{X}_j)} l^*(\mathbf{X}_i, \tilde{\mathbf{x}}), \quad (1.3)$$

over the product space $\{1, \dots, N\} \times \tilde{\Xi}$ of indices and particle positions, where we have introduced the partition function $a^*(\mathbf{x}) := \int l(\mathbf{x}, \tilde{\mathbf{x}}) d\tilde{\mathbf{x}}$ of $l(\mathbf{x}, \cdot)$ and the normalized (Markovian) transition density $l^*(\mathbf{x}, \tilde{\mathbf{x}}) := l(\mathbf{x}, \tilde{\mathbf{x}}) / \int l(\mathbf{x}, \tilde{\mathbf{x}}) d\tilde{\mathbf{x}}$. To sample (1.3), we take an importance sampling approach and draw independent pairs $\{(I_i, \tilde{\mathbf{X}}_i)\}_{i=1}^N$ of indices and positions from the proposal distribution

$$\pi(i, \tilde{\mathbf{x}}) := \frac{\omega_i a(\mathbf{X}_i)}{\sum_{j=1}^N \omega_j a(\mathbf{X}_j)} r(\mathbf{X}_i, \tilde{\mathbf{x}}) \quad (1.4)$$

over the same extended space and assign each draw $(I_i, \tilde{\mathbf{X}}_i)$ the importance weight $\tilde{\omega}_i := w(I_i, \tilde{\mathbf{X}}_i)$, where

$$w(i, \tilde{\mathbf{x}}) := \frac{l(\mathbf{X}_i, \tilde{\mathbf{x}})}{a(\mathbf{X}_i) r(\mathbf{X}_i, \tilde{\mathbf{x}})} \propto \frac{\mu^{\text{aux}}(i, \tilde{\mathbf{x}})}{\pi(i, \tilde{\mathbf{x}})}.$$

Finally, we discard the indices I_i and keep $\{(\tilde{\mathbf{X}}_i, \tilde{\omega}_i)\}_{i=1}^N$ as an approximation of μ . This scheme, which is traditionally referred to as the *auxiliary particle filter*, encompasses the simpler framework of the *bootstrap particle filter* proposed originally by Gordon et al. (1993), which simply amounts to setting $a(\mathbf{x}) = 1$ for all \mathbf{x} (implying a gain of simplicity at the price of flexibility). Moreover, SMC schemes where resampling is performed only at random times can similarly be cast into the setting of the auxiliary particle filter by composing the kernels involved in several consecutive steps of the of recursion (1.1) (see e.g. Cornebise, 2009, Chapter 4, for details). Therefore, any methodology built for the auxiliary particle filter also applies to these simpler frameworks without modification.

The choice of the nonnegative *adjustment multiplier function* a and the *proposal kernel* r (which is supposed to dominate l^*) effects significantly the quality of the returned sample, and in this paper we focus on adaptive design of the latter. Doucet et al. (2000) suggest to approximate each function $l(\mathbf{X}_i, \cdot)$ by a Gaussian density whose mean and covariance are obtained using the extended Kalman filter (EKF). This technique is computationally heavy for large sample sizes since it requires to solve an optimization problem for each particle \mathbf{x}_i . A somewhat simpler version, consisting in replacing the EKF by the unscented Kalman filter (UKF), was studied later by Van der Merwe et al. (2000); Van der Merwe and Wan (2003). Pitt and Shephard (1999) use a form of Laplace approximations for approximating $l(\mathbf{X}_i, \cdot)$ by simply centering a Gaussian density or the density of a student's t -distribution around the mode of the function in question. This technique is appropriate when the function is log-concave (or strongly unimodal).

Nevertheless, the localization of the mode requires solving an optimization problem for each particle. Thus, in spite of this intense recent activity in the field, the state-of-the-art algorithms have met only mitigated success as they implicitly assume that the functions $l(\mathbf{X}_i, \cdot)$ have a single mode.

A common practice (see e.g. [Oh and Berger, 1993](#)) for designing proposal distributions in standard (non-sequential) importance sampling is to consider a parametric family of proposal distributions and then identify a parameter that minimizes some measure of discrepancy between the target and the proposal distributions. Common choices are the widely used squared *coefficient of variation* (CV)

$$\text{CV}^2 := N \sum_{i=1}^N \left(\frac{\tilde{\omega}_i}{\bar{\tilde{\omega}}} \right)^2 - 1$$

of the importance weights or the negated *Shannon entropy*

$$\mathcal{E} := \sum_{i=1}^N \frac{\tilde{\omega}_i}{\bar{\tilde{\omega}}} \log \left(\frac{N \tilde{\omega}_i}{\bar{\tilde{\omega}}} \right)$$

of the same. Both are maximal when the sample is completely degenerated, i.e. when one particle carries all the weight, and minimal when all the importance weights are equal. Note that the coefficient of variation often appears in a transformed form as the *effective sample size* given by $\text{ESS} = N/(1 + \text{CV}^2)$.

The same quantities have been widely used also within the framework of SMC, but until recently only for triggering the resampling operation and not as a tool for adaptive design of the instrumental distribution. The key result of [Cornebise et al. \(2008\)](#) was to relate the Shannon entropy to the Kullback-Leibler divergence (KLD, recalled below) between the instrumental and proposal distributions of the particle filter. In addition, a similar relation between the CV of the particle weights and the Chi-square divergence (CSD) between the same distributions was established. More specifically, as the number of particles tend to infinity, the CV and the entropy tend to the CSD resp. the KLD between the asymptotic target distribution $\mu^*(\mathbf{x}, \tilde{\mathbf{x}}) \propto l(\mathbf{x}, \tilde{\mathbf{x}})\nu(\mathbf{x})$ and the asymptotic proposal distribution $\pi^*(\mathbf{x}, \tilde{\mathbf{x}}) \propto a(\mathbf{x})r(\mathbf{x}, \tilde{\mathbf{x}})\nu(\mathbf{x})$ on the product space $\mathbf{X} \times \tilde{\mathbf{X}}$, i.e. under weak integrability hypotheses, as $N \rightarrow \infty$,

$$\begin{aligned} \mathcal{E} &\xrightarrow{\mathbb{P}} d_{\text{KL}}(\mu^* \| \pi^*) , & \text{CV}^2 &\xrightarrow{\mathbb{P}} d_{\chi^2}(\mu^* \| \pi^*) , \\ d_{\text{KL}}(\mu^{\text{aux}} \| \pi) &\xrightarrow{\mathbb{P}} d_{\text{KL}}(\mu^* \| \pi^*) , & d_{\chi^2}(\mu^{\text{aux}} \| \pi) &\xrightarrow{\mathbb{P}} d_{\chi^2}(\mu^* \| \pi^*) . \end{aligned}$$

This gives a sound theoretical support for using the CV and the entropy of the importance weights for measuring the quality of the particle sample. In addition, it suggests that the KLD and the CSD between μ^{aux} and π could be used in lieu of \mathcal{E} and CV^2 , respectively, for all purposes, especially adaptation. As a matter of fact, in the context of adaptive design of SMC methods, the KLD is (as pointed

out in [Cornebise et al., 2008](#), Section 2.3) highly practical since it decouples the adaptation of the adjustment weight function a and that of the proposal kernel r . Recall that the KLD between two distributions with densities p and q over some space Ξ is defined as

$$d_{\text{KL}}(p||q) := \int_{\Xi} \log \left(\frac{p(\mathbf{x})}{q(\mathbf{x})} \right) p(\mathbf{x}) d\mathbf{x} ,$$

provided that the support of p is included in the support of q . Replacing p and q by the auxiliary target distribution μ^{aux} and the proposal π , respectively, yields

$$\begin{aligned} d_{\text{KL}}(\mu^{\text{aux}}||\pi) &= \mathbb{E}_{\mu^{\text{aux}}} \left[\log \frac{\mu^{\text{aux}}(I, \tilde{\mathbf{X}})}{\pi(I, \tilde{\mathbf{X}})} \right] \\ &= \underbrace{\mathbb{E}_{\mu^{\text{aux}}} \left[\log \frac{l^*(\mathbf{X}_I, \tilde{\mathbf{X}})}{r(\mathbf{X}_I, \tilde{\mathbf{X}})} \right]}_{\text{Depends only on } r} + \underbrace{\mathbb{E}_{\mu^{\text{aux}}} \left[\log \frac{\omega_I a^*(\mathbf{X}_I) / \sum_{j=1}^N \omega_j a^*(\mathbf{X}_j)}{\omega_I a(\mathbf{X}_I) / \sum_{j=1}^N \omega_j a(\mathbf{X}_j)} \right]}_{\text{Depends only on } a} . \end{aligned} \quad (1.5)$$

The first term corresponds to the discrepancy induced by proposing the particles $\tilde{\mathbf{X}}$ with a suboptimal proposal kernel, and the second term corresponds to the discrepancy induced by choosing the ancestor indices I according to suboptimal adjustment weights. Moreover, $d_{\text{KL}}(\mu^{\text{aux}}||\pi)$ equals

$$-\mathbb{E}_{\mu^{\text{aux}}} \left[\log r(\mathbf{X}_I, \tilde{\mathbf{X}}) \right] - \mathbb{E}_{\mu^{\text{aux}}} \left[\log \frac{a(\mathbf{X}_I)}{\sum_{j=1}^N \omega_j a(\mathbf{X}_j)} \right] \quad (1.7)$$

up to additive terms involving the optimal quantities a^* and l^* only and thus being irrelevant for the optimization problem (equality up to a constant will in the following be denoted by \equiv). Restricting ourselves to the adaptation of the proposal kernel, we obtain the simple expression $d_{\text{KL}}(\mu^{\text{aux}}||\pi) \equiv -\mathbb{E}_{\mu^{\text{aux}}} \left[\log r(\mathbf{X}_I, \tilde{\mathbf{X}}) \right]$. Although the d_{KL} is most often not directly computable, such a quantity can be approximated *on-the-fly* using the weighted sample obtained; the resulting algorithm closely resembles the cross-entropy (CE; see [Rubinstein and Kroese, 2004](#)) method.

Henceforth, in the present article, we select the auxiliary proposal distribution π from a family $\{\pi_{\boldsymbol{\theta}} : \boldsymbol{\theta} \in \Theta\}$ of candidates by first solving $\boldsymbol{\theta}^* := \arg \min_{\boldsymbol{\theta} \in \Theta} d_{\text{KL}}(\mu^{\text{aux}}||\pi_{\boldsymbol{\theta}})$ and then letting $\pi = \pi_{\boldsymbol{\theta}^*}$. On the one hand, the chosen parametric family should be flexible enough to approximate complex transition kernels. On the other hand, it should be simple enough so that sampling from $\pi_{\boldsymbol{\theta}}$ is easy. Finally, the parameterization should be done in such a way that the problem of estimating the parameters is as simple as possible. In this article, we suggest modeling the proposal $\pi_{\boldsymbol{\theta}}$ as a mixture of integrated curved exponential

distributions. The mixture weights are chosen to be dependent on the ancestor particle \mathbf{x}_j , allowing for partitioning of the input space into regions corresponding to a specialized kernel. Each component of the mixture belongs to a family of integrated curved exponential distributions, whose two most known members are the multivariate Gaussian distribution and the student's t -distribution. Also the parameters of the mixture depend on the particle positions \mathbf{x}_j . This parameterization of the proposal distribution is closely related to the (hierarchical) *mixture of experts* appearing in the machine learning community and described in (Jordan and Jacobs, 1994; Jordan and Xu, 1995). The flexibility of our approach allows for efficient approximation of the optimal kernel for a large class of intricate (nonlinear/non-Gaussian) models.

The paper is organized as follows. The mixture family under consideration is described in Section 2 and in Section 3 we treat optimization of the parameters. In Section 4, two versions of the adaptive particle filters are presented and in the last part, Section 5, we illustrate the efficiency of the method on several simulation-based examples.

2 Mixture of experts

In this contribution, we let the proposal kernel have density

$$r_{\boldsymbol{\theta}}(\mathbf{x}, \tilde{\mathbf{x}}) := \sum_{j=1}^d \alpha_j(\mathbf{x}, \underline{\boldsymbol{\beta}}) \rho(\mathbf{x}, \tilde{\mathbf{x}}; \boldsymbol{\eta}_j) , \quad (2.1)$$

where the functions $\{\alpha_j\}_{j=1}^d$ are so-called *weighting functions* and $\rho(\cdot, \cdot; \boldsymbol{\eta}_j)$ are Markov transition kernels from Ξ to $\tilde{\Xi}$. Denote $\boldsymbol{\theta} := (\underline{\boldsymbol{\beta}}, \underline{\boldsymbol{\eta}})$ with $\underline{\boldsymbol{\eta}} = (\boldsymbol{\eta}_1, \dots, \boldsymbol{\eta}_d)$. In the auxiliary particle filter framework, the associated proposal distribution is defined as

$$\pi_{\boldsymbol{\theta}}(i, \tilde{\mathbf{x}}) := \frac{\omega_i a(\mathbf{X}_i)}{\sum_{k=1}^N \omega_k a(\mathbf{X}_k)} \sum_{j=1}^d \alpha_j(\mathbf{X}_i, \underline{\boldsymbol{\beta}}) \rho(\mathbf{X}_i, \tilde{\mathbf{x}}; \boldsymbol{\eta}_j) . \quad (2.2)$$

We then assign the importance weight $\tilde{w}_i^{\boldsymbol{\theta}} = w^{\boldsymbol{\theta}}(I_i, \tilde{\mathbf{X}}_i)$ to each draw $(I_i, \tilde{\mathbf{X}}_i)$ from $\pi_{\boldsymbol{\theta}}$, where

$$w^{\boldsymbol{\theta}}(i, \tilde{\mathbf{x}}) = \frac{l(\mathbf{X}_i, \tilde{\mathbf{x}})}{a(\mathbf{X}_i) \sum_{j=1}^d \alpha_j(\mathbf{X}_i, \underline{\boldsymbol{\beta}}) \rho(\mathbf{X}_i, \tilde{\mathbf{x}}; \boldsymbol{\eta}_j)} \propto \frac{\mu^{\text{aux}}(i, \tilde{\mathbf{x}})}{\pi_{\boldsymbol{\theta}}(i, \tilde{\mathbf{x}})} . \quad (2.3)$$

The weighting function partitions the input space Ξ into regions which are associated to a few *specialized* transition kernels in the mixture. This is done by assigning a vector of mixture weights to each point of the input space. As in Jordan and Xu (1995), we consider logistic weight functions, i.e. $\underline{\boldsymbol{\beta}} = (\boldsymbol{\beta}_1, \dots, \boldsymbol{\beta}_{d-1})$

and

$$\alpha_j(\mathbf{x}, \underline{\boldsymbol{\beta}}) := \frac{\exp(\boldsymbol{\beta}_j^T \bar{\mathbf{x}})}{1 + \sum_{k=1}^{d-1} \exp(\boldsymbol{\beta}_k^T \bar{\mathbf{x}})}, \quad 1 \leq j \leq d-1, \quad (2.4)$$

$$\alpha_d(\mathbf{x}, \underline{\boldsymbol{\beta}}) := \frac{1}{1 + \sum_{k=1}^{d-1} \exp(\boldsymbol{\beta}_k^T \bar{\mathbf{x}})}, \quad (2.5)$$

where $\bar{\mathbf{x}} := (\mathbf{x}, 1)$ and $\boldsymbol{\beta}_j$, $1 \leq j \leq d-1$, are vectors in \mathbb{R}^{p+1} .

In some cases it is of interest to resort to a simpler mixtures whose weights do not depend on \mathbf{x} by letting, for $\underline{\boldsymbol{\beta}} \in \mathbb{R}_+^d$ such that $\sum_{j=1}^d \boldsymbol{\beta}_j = 1$, $\alpha_j(\mathbf{x}, \underline{\boldsymbol{\beta}}) := \boldsymbol{\beta}_j$, $1 \leq j \leq d$, independently of \mathbf{x} and hence without partitioning the input space. This model for the transition kernel is then similar to the *switching regression model* (Quandt and Ramsey, 1972).

To ease the implementation—see Section 3—the kernels $\rho(\mathbf{x}, \tilde{\mathbf{x}}; \boldsymbol{\eta})$ are assumed to have the following *integrated curved exponential* form:

$$\rho(\mathbf{x}, \tilde{\mathbf{x}}; \boldsymbol{\eta}) = \int_{\mathbb{U}} \rho^e(\mathbf{x}, \tilde{\mathbf{x}}, \mathbf{u}; \boldsymbol{\eta}) d\mathbf{u}, \quad (2.6)$$

where $\rho^e(\mathbf{x}, \tilde{\mathbf{x}}, \mathbf{u}; \boldsymbol{\eta})$ is a *curved exponential distribution*

$$\rho^e(\mathbf{x}, \tilde{\mathbf{x}}, \mathbf{u}; \boldsymbol{\eta}) = \gamma(\mathbf{u}) h(\mathbf{x}, \tilde{\mathbf{x}}, \mathbf{u}) \exp\left(-A(\boldsymbol{\eta}) + \text{Tr}\left(S(\mathbf{x}, \tilde{\mathbf{x}}, \mathbf{u})^T B(\boldsymbol{\eta})\right)\right). \quad (2.7)$$

In other words, $\rho(\mathbf{x}, \tilde{\mathbf{x}}; \boldsymbol{\eta})$ corresponds to the marginal (in $(\mathbf{x}, \tilde{\mathbf{x}})$) of a curved exponential distribution $\rho^e(\mathbf{x}, \tilde{\mathbf{x}}, \mathbf{u}; \boldsymbol{\eta})$, where \mathbf{u} is an auxiliary variable taking values in some space \mathbb{U} and γ is the marginal distribution of \mathbf{u} . In (2.7), S is the vector of sufficient statistics.

We now augment $(I, \tilde{\mathbf{X}})$ with the index J of the mixture component as well as the auxiliary variable \mathbf{U} of the curved exponential family. This extended auxiliary variable $(I, J, \tilde{\mathbf{X}}, \mathbf{U})$ is distributed according to

$$\pi_{\boldsymbol{\theta}}^e(i, j, \tilde{\mathbf{x}}, \mathbf{u}) := \frac{\omega_i a(\mathbf{X}_i)}{\sum_{i=1}^N \omega_i a(\mathbf{X}_i)} \alpha_j(\mathbf{X}_i, \underline{\boldsymbol{\beta}}) \rho^e(\mathbf{X}_i, \tilde{\mathbf{x}}, \mathbf{u}; \boldsymbol{\eta}_j) \quad (2.8)$$

$$\pi_{\boldsymbol{\theta}}^e(i, j, \tilde{\mathbf{x}}, \mathbf{u}) := \frac{\omega_i a(\mathbf{X}_i)}{\sum_{k=1}^N \omega_k a(\mathbf{X}_k)} \alpha_j(\mathbf{X}_i, \underline{\boldsymbol{\beta}}) \rho^e(\mathbf{X}_i, \tilde{\mathbf{x}}, \mathbf{u}; \boldsymbol{\eta}_j)$$

on the product space $\{1, \dots, N\} \times \{1, \dots, d\} \times \tilde{\Xi} \times \mathbb{U}$. It is easily checked that $\pi_{\boldsymbol{\theta}}$ is the marginal density of $\pi_{\boldsymbol{\theta}}^e$ in I and $\tilde{\mathbf{X}}$.

Efficient fitting of the optimal parameter $\boldsymbol{\theta}^*$ minimizing the KLD (1.7) is the topic of the Section 3. We will illustrate the details for the specific case of multidimensional Gaussian distributions and multidimensional t -distributions. It should however be kept in mind that the algorithm is valid for any member of the family of integrated curved exponential distributions as long as Assumption 3.1 below is fulfilled.

3 Parameter estimation techniques

Parameter estimation for mixtures of experts is most often carried through using the EM algorithm (see [McLachlan and Krishnan, 2008](#), Section 8.4.6) and also in this paper a recursive EM-type algorithm will be used for finding optimal mixture parameters. In order to describe precisely the algorithm, let $\boldsymbol{\theta}^{\ell-1}$ be a given parameter estimate, obtained after $\ell - 1$ iterations of the scheme. At iteration ℓ , a new fit $\boldsymbol{\theta}^\ell$ of the parameter is obtained as the argument $\boldsymbol{\theta}$ maximizing the *intermediate quantity* $Q(\boldsymbol{\theta}, \boldsymbol{\theta}^{\ell-1}) := \mathbb{E}_{\mu^{\text{aux}}} [\mathbb{E}_{\boldsymbol{\theta}^{\ell-1}} [\log \pi_{\boldsymbol{\theta}}^e(I, J, \tilde{\mathbf{X}}, \mathbf{U}) | I, \tilde{\mathbf{X}}]]$, where $\pi_{\boldsymbol{\theta}}^e$ is the extended auxiliary distribution in (2.8) and $\mathbb{E}_{\boldsymbol{\theta}^{\ell-1}}$ denotes expectation w.r.t. the distribution $\pi_{\boldsymbol{\theta}^{\ell-1}}^e$. The intermediate quantity may be decomposed as $Q(\boldsymbol{\theta}, \boldsymbol{\theta}^{\ell-1}) \equiv Q_1(\underline{\boldsymbol{\beta}}, \boldsymbol{\theta}^{\ell-1}) + Q_2(\underline{\boldsymbol{\eta}}, \boldsymbol{\theta}^{\ell-1})$, where

$$Q_1(\underline{\boldsymbol{\beta}}, \boldsymbol{\theta}^{\ell-1}) := \mathbb{E}_{\mu^{\text{aux}}} \left[\mathbb{E}_{\boldsymbol{\theta}^{\ell-1}} \left[\log \alpha_J(\mathbf{X}_I, \underline{\boldsymbol{\beta}}) \middle| I, \tilde{\mathbf{X}} \right] \right], \quad (3.1)$$

$$Q_2(\underline{\boldsymbol{\eta}}, \boldsymbol{\theta}^{\ell-1}) := \mathbb{E}_{\mu^{\text{aux}}} \left[\mathbb{E}_{\boldsymbol{\theta}^{\ell-1}} \left[\log \rho^e(\mathbf{X}_I, \tilde{\mathbf{X}}, \mathbf{U}; \underline{\boldsymbol{\eta}}_J) \middle| I, \tilde{\mathbf{X}} \right] \right]. \quad (3.2)$$

Therefore, at each iteration, maximization of the intermediate quantity may be split into the two subproblems $\underline{\boldsymbol{\beta}}^\ell = \arg \max_{\underline{\boldsymbol{\beta}}} Q_1(\underline{\boldsymbol{\beta}}, \boldsymbol{\theta}^{\ell-1})$ and $\underline{\boldsymbol{\eta}}^\ell = \arg \max_{\underline{\boldsymbol{\eta}}} Q_2(\underline{\boldsymbol{\eta}}, \boldsymbol{\theta}^{\ell-1})$. Under weak additional assumptions, it suffices to find a value $\boldsymbol{\theta}$ increasing $Q(\boldsymbol{\theta}, \boldsymbol{\theta}^{\ell-1})$ over its baseline, i.e. $Q(\boldsymbol{\theta}^\ell, \boldsymbol{\theta}^{\ell-1}) > Q(\boldsymbol{\theta}^{\ell-1}, \boldsymbol{\theta}^{\ell-1})$, in order to obtain a sequence $\{\boldsymbol{\theta}^\ell\}_{\ell \geq 0}$ converging to the optimum $\boldsymbol{\theta}^*$; this is of special importance when using logistic weights, as finding the exact optimum of the intermediate quantity could be expensive in that case.

3.1 Updating of logistic weight parameters

For the logistic weights (2.4), finding the maximum in $\underline{\boldsymbol{\beta}}$ of $Q_1(\underline{\boldsymbol{\beta}}, \boldsymbol{\theta}^\ell)$ on closed form is out of reach but can be performed using *iteratively reweighted least squares* (IRLS), as in ([McLachlan and Krishnan, 2008](#), Equations (8.32) and (8.33)).

By defining the *conditional mixture weights*

$$\pi_{\boldsymbol{\theta}}^e(j|i, \tilde{\mathbf{x}}) = \frac{\alpha_j(\mathbf{X}_i, \underline{\boldsymbol{\beta}}) \rho(\mathbf{X}_i, \tilde{\mathbf{x}}; \boldsymbol{\eta}_j)}{\sum_{k=1}^d \alpha_k(\mathbf{X}_i, \underline{\boldsymbol{\beta}}) \rho(\mathbf{X}_i, \tilde{\mathbf{x}}; \boldsymbol{\eta}_k)}, \quad (3.3)$$

we may express the k th block (of size $(p+1)$) of the gradient vector $\nabla_{\underline{\boldsymbol{\beta}}} Q_1(\underline{\boldsymbol{\beta}}, \boldsymbol{\theta}^{\ell-1})$ as

$$\nabla_{\boldsymbol{\beta}_k} Q_1(\underline{\boldsymbol{\beta}}, \boldsymbol{\theta}^{\ell-1}) = \mathbb{E}_{\mu^{\text{aux}}} \left[\left\{ \pi_{\boldsymbol{\theta}^{\ell-1}}^e \left(k \middle| I, \tilde{\mathbf{X}} \right) - \alpha_k(\mathbf{X}_I, \underline{\boldsymbol{\beta}}) \right\} \bar{\mathbf{X}}_I \right] \quad (3.4)$$

for any $k \in \{1, \dots, d-1\}$. Similarly, for any $(k, m) \in \{1, \dots, d-1\}^2$, block (k, m)

(of size $(p+1) \times (p+1)$) of the Hessian matrix $\nabla_{\underline{\beta}}^2 Q_1(\underline{\beta}, \boldsymbol{\theta}^\ell)$ can be expressed as

$$\nabla_{\beta_k} \nabla_{\beta_m} Q_1(\underline{\beta}, \boldsymbol{\theta}^{\ell-1}) = \mathbb{E}_{\mu^{\text{aux}}} \left[\alpha_k(\mathbf{X}_I, \underline{\beta}) \left\{ \alpha_m(\mathbf{X}_I, \underline{\beta}) - \frac{\mathbb{1}_{\{k=m\}}}{1 + \sum_{l=1}^{d-1} \exp(\beta_l^T \bar{\mathbf{X}}_I)} \right\} \bar{\mathbf{X}}_I^2 \right], \quad (3.5)$$

where, for any vector or matrix A , A^2 is shorthand for $A A^T$. Note that the Hessian is the conditional expectation of the complete data negated observed Fisher information matrix, and is therefore definite negative. These results stem directly from the multinomial logistic regression by computing conditional expectations. The EM intermediate quantity $\underline{\beta} \mapsto Q_1(\underline{\beta}, \boldsymbol{\theta}^{\ell-1})$ is therefore concave. We could thus achieve exact optimization at each iteration ℓ using IRLS or running the Newton-Raphson scheme until convergence (Chen et al., 1999, who also correct the erroneous expression of the Hessian matrix in IRLS), which is commonly done in multinomial logistic regression. However, a single Newton-Raphson step is enough to ensure that the intermediate quantity is increased, which implies convergence of the EM scheme. Consequently, an EM iteration may be written as

$$\underline{\beta}^\ell := \underline{\beta}^{\ell-1} + \tau_\ell \left[\nabla_{\underline{\beta}}^2 Q_1(\underline{\beta}, \boldsymbol{\theta}^{\ell-1})|_{\underline{\beta}=\underline{\beta}^{\ell-1}} \right]^{-1} \nabla_{\underline{\beta}} Q_1(\underline{\beta}, \boldsymbol{\theta}^{\ell-1})|_{\underline{\beta}=\underline{\beta}^{\ell-1}} \quad (3.6)$$

for some positive step size (or *learning rate*) $\tau_\ell \leq 1$. Following McLachlan and Krishnan (2008), the learning rate is set to $\tau_\ell = 1$.

3.2 Updating of strata parameters

We now turn to estimation of the strata parameters $\underline{\eta}$.

Assumption 3.1 (Closed form of integrated sufficient statistics). In the sequel, it is assumed that, for any $\underline{\eta}$ and any $j \in \{1, \dots, d\}$, the *integrated sufficient statistics*

$$\int S(\mathbf{X}_i, \tilde{\mathbf{x}}, \mathbf{u}) \rho^e(\mathbf{u}|\mathbf{X}_i, \tilde{\mathbf{x}}; \boldsymbol{\eta}_j) d\mathbf{u} \quad (3.7)$$

are available on closed form. This is for instance the case for Student's t -distributions.

Under Assumption 3.1, for any $j \in \{1, \dots, d\}$ we define the *expected sufficient statistics of component j* as

$$P_j(\boldsymbol{\theta}) := \mathbb{E}_{\mu^{\text{aux}}} \left[\pi_{\boldsymbol{\theta}}^e(j|I, \tilde{\mathbf{X}}) \right], \quad (3.8)$$

$$S_j(\boldsymbol{\theta}) := \mathbb{E}_{\mu^{\text{aux}}} \left[\pi_{\boldsymbol{\theta}}^e(j|I, \tilde{\mathbf{X}}) \int S(\mathbf{X}_I, \tilde{\mathbf{X}}, \mathbf{u}) \rho^e(\mathbf{u}|\mathbf{X}_I, \tilde{\mathbf{X}}; \boldsymbol{\eta}_j) d\mathbf{u} \right]. \quad (3.9)$$

We now derive the explicit update expressions for the two common special cases of multivariate Gaussian and multivariate Student's t -distributions; it should however be kept in mind that the method is valid for any integrated curved exponential distribution satisfying Assumption 3.1.

3.2.1 Updating formulas for Gaussian strata

To each component j , assign a linear Gaussian regression parameterized by $\boldsymbol{\eta}_j = (M_j, \Sigma_j)$. Thus, given $J = j$, the new particle $\tilde{\mathbf{X}}$ has p' -dimensional Gaussian distribution with mean $M_j \bar{\mathbf{X}}_I$, where the *regression matrix* M_j is of size $p' \times (p+1)$, and symmetric, positive definite *covariance matrix* Σ_j of size $p' \times p'$. In this case, the sufficient statistics are $S(\mathbf{x}, \tilde{\mathbf{x}}, \mathbf{u}) = (\tilde{\mathbf{x}}^2, \bar{\mathbf{x}}^2, \tilde{\mathbf{x}} \bar{\mathbf{x}}^T)$, leading to the expected sufficient statistics

$$\begin{aligned} P_j(\boldsymbol{\theta}^{\ell-1}) &:= \mathbb{E}_{\mu^{\text{aux}}} \left[\pi_{\boldsymbol{\theta}^{\ell-1}}^e \left(j \middle| I, \tilde{\mathbf{X}} \right) \right], & S_{j,1}(\boldsymbol{\theta}^{\ell-1}) &:= \mathbb{E}_{\mu^{\text{aux}}} \left[\pi_{\boldsymbol{\theta}^{\ell-1}}^e \left(j \middle| I, \tilde{\mathbf{X}} \right) \tilde{\mathbf{X}}^2 \right], \\ S_{j,2}(\boldsymbol{\theta}^{\ell-1}) &:= \mathbb{E}_{\mu^{\text{aux}}} \left[\pi_{\boldsymbol{\theta}^{\ell-1}}^e \left(j \middle| I, \tilde{\mathbf{X}} \right) \bar{\mathbf{X}}_I^2 \right], & S_{j,3}(\boldsymbol{\theta}^{\ell-1}) &:= \mathbb{E}_{\mu^{\text{aux}}} \left[\pi_{\boldsymbol{\theta}^{\ell-1}}^e \left(j \middle| I, \tilde{\mathbf{X}} \right) \tilde{\mathbf{X}} \bar{\mathbf{X}}_I^T \right]. \end{aligned} \quad (3.10)$$

In addition, the functions A and B in (2.7) are given by

$$A(\boldsymbol{\eta}) = \frac{1}{2} \log |\Sigma|, \quad B(\boldsymbol{\eta}) = \left(-\frac{1}{2} \Sigma^{-1}, -\frac{1}{2} M \Sigma^{-1} M^T, -\Sigma^{-1} M \right).$$

Consequently, the argument $\underline{\boldsymbol{\eta}}$ maximizing $Q_2(\underline{\boldsymbol{\eta}}, \boldsymbol{\theta}^{\ell-1})$ is given by $M_j^\ell = M(S_j(\boldsymbol{\theta}^{\ell-1}))$ and $\Sigma_j^\ell = \Sigma(P_j(\boldsymbol{\theta}^{\ell-1}), S_j(\boldsymbol{\theta}^{\ell-1}))$, where the functions M and Σ are defined by

$$M(S_j) := S_{j,3} S_{j,2}^{-1} \quad \text{and} \quad \Sigma(P_j, S_j) := \frac{S_{j,1} - S_{j,3} S_{j,2}^{-1} S_{j,3}^T}{P_j}. \quad (3.11)$$

Remark 3.1 (Pooling the covariances). In practice, a robustified version of the algorithm above can be obtained by *pooling* the covariances. This means that a common covariance matrix is used for all the components of the mixture, i.e. $\Sigma_j = \Sigma$ for every j . By doing this, the well-known problem of mixture models with strata components degenerating to Dirac masses is avoided. It is straightforward to enforce such a restriction in the optimization procedure above, leading to

$$\Sigma_j^\ell = \Sigma^\ell := \frac{1}{d} \sum_{k=1}^d (S_{k,1} - S_{k,3} S_{k,2}^{-1} S_{k,3}^T) \quad (3.12)$$

for all j .

3.2.2 Updating formulas for Student's t -distributed strata

A common advice in Importance Sampling (see, for instance, [Oh and Berger, 1993](#)) is to replace Gaussian distributions by Student's t -distributions with location parameter μ , scale parameter Σ , and a chosen number $\nu > 0$ of degrees of freedom. Therefore, as an alternative, robustified version of the Gaussian mixture model above, we propose using instead a p' -dimensional Student's t -distribution with ν degrees of freedom, i.e. $\rho(\mathbf{x}, \tilde{\mathbf{x}}; \boldsymbol{\eta}_j) = t_{p'}(\tilde{\mathbf{x}}; M_j \bar{\mathbf{x}}, \Sigma_j, \nu)$.

Remark 3.2 (Fixing the degrees of freedom). The number ν of degrees of freedom of the Student's t -distributions is fixed, typically to $\nu \in \{3, 4\}$, beforehand and is common to all the components. A similar choice has been made by, among others, [Peel and McLachlan \(2000, Section 7\)](#) and [Cappé et al. \(2008\)](#) as it allows for closed form optimization in the M -step.

These kernels can be cast into the framework of [Section 2](#) thanks to the *Gaussian-Gamma decomposition* of multivariate Student's t -distributions used in ([Liu and Rubin, 1995, Section 2](#)) and ([Peel and McLachlan, 2000, Section 3](#)),

$$t(\tilde{\mathbf{x}}; \mu, \Sigma, \nu) = \int \mathcal{N}(\tilde{\mathbf{x}}; \mu, \Sigma/\mathbf{u}) \gamma\left(\mathbf{u}; \frac{\nu}{2}, \frac{\nu}{2}\right) d\mathbf{u},$$

where $\gamma(\mathbf{u}; a, b) := b^a \mathbf{u}^{a-1} \exp(-b\mathbf{u}) \mathbb{1}_{\mathbb{R}^+}(\mathbf{u})/\Gamma(a)$ is the density of the Gamma distribution with shape parameter a and scale parameter b . The multivariate Student's t -distribution kernel is hence an integrated curved exponential distribution [\(2.6\)](#) with $\gamma(\mathbf{u}) = \gamma(\mathbf{u}; \frac{\nu}{2}, \frac{\nu}{2})$, $h(\mathbf{x}, \tilde{\mathbf{x}}, \mathbf{u}) = (2\pi)^{-p/2}$, sufficient statistics $S(\mathbf{x}, \tilde{\mathbf{x}}, \mathbf{u}) = (\mathbf{u}\tilde{\mathbf{x}}^2, \mathbf{u}\bar{\mathbf{x}}^2, \mathbf{u}\tilde{\mathbf{x}}\bar{\mathbf{x}}^T)$, and the same A and B as in the Gaussian case. Since the Gamma distribution is a conjugate prior w.r.t. precision for the Gaussian distribution with known mean, it holds that

$$\mathbb{E}_{\boldsymbol{\eta}}[\mathbf{U}|\mathbf{x}, \tilde{\mathbf{x}}] = \frac{\nu + p'}{\nu + \delta(\tilde{\mathbf{x}}, M\bar{\mathbf{x}}, \Sigma)}, \quad (3.13)$$

where $\delta(\tilde{\mathbf{x}}, \mu, \Sigma) = |\tilde{\mathbf{x}} - \mu|_{\Sigma^{-1}}^2 = (\tilde{\mathbf{x}} - \mu)^T \Sigma^{-1} (\tilde{\mathbf{x}} - \mu)$ is the Mahalanobis distance with covariance matrix Σ . This leads to the expected sufficient statistics

$$\begin{aligned} P_j(\boldsymbol{\theta}^{\ell-1}) &:= \mathbb{E}_{\mu^{\text{aux}}} \left[\pi_{\boldsymbol{\theta}^{\ell-1}}^e(j|I, \tilde{\mathbf{X}}) \right], \\ S_{j,1}(\boldsymbol{\theta}^{\ell-1}) &:= \mathbb{E}_{\mu^{\text{aux}}} \left[\pi_{\boldsymbol{\theta}^{\ell-1}}^e(j|I, \tilde{\mathbf{X}}) \frac{\nu + p'}{\nu + \delta(\tilde{\mathbf{X}}, M_j \bar{\mathbf{X}}_I, \Sigma_j)} \tilde{\mathbf{X}}^2 \right], \\ S_{j,2}(\boldsymbol{\theta}^{\ell-1}) &:= \mathbb{E}_{\mu^{\text{aux}}} \left[\pi_{\boldsymbol{\theta}^{\ell-1}}^e(j|I, \tilde{\mathbf{X}}) \frac{\nu + p'}{\nu + \delta(\tilde{\mathbf{X}}, M_j \bar{\mathbf{X}}_I, \Sigma_j)} \bar{\mathbf{X}}_I^2 \right], \\ S_{j,3}(\boldsymbol{\theta}^{\ell-1}) &:= \mathbb{E}_{\mu^{\text{aux}}} \left[\pi_{\boldsymbol{\theta}^{\ell-1}}^e(j|I, \tilde{\mathbf{X}}) \frac{\nu + p'}{\nu + \delta(\tilde{\mathbf{X}}, M_j \bar{\mathbf{X}}_I, \Sigma_j)} \tilde{\mathbf{X}} \bar{\mathbf{X}}_I^T \right]. \end{aligned} \quad (3.14)$$

Since the functions A and B are unchanged from the Gaussian case, the same equations [\(3.11\)](#) can be applied for updating the parameter $\boldsymbol{\eta}$ based on the expected sufficient statistics [\(3.14\)](#). Covariance pooling is also achieved in exactly same way as in [Remark 3.1](#).

4 Stochastic approximation and resulting algorithm

The algorithms described so far are mainly of theoretical interest, since computation of the sufficient statistics $S_{j,i}(\boldsymbol{\theta})$ defined in (3.9) involves computing expectations w.r.t. to the target distribution μ^{aux} . We may however approximate the same using importance sampling with a relatively small sample size. This is described in the following.

At iteration ℓ , we sample a block of N^ℓ particles from $\pi_{\boldsymbol{\theta}^{\ell-1}}$ and form each estimate $\hat{S}_{j,i}^\ell$ of an expected sufficient statistic $S_{j,i}$ as a convex combination of the former estimate $\hat{S}_{j,i}^{\ell-1}$ and an importance sampling approximation based on the obtained particle block. The weights of the convex combination are given by a sequence $\{\lambda_\ell\}_{\ell=0}^\infty$ of positive step sizes such that $\sum_{\ell=0}^\infty \lambda_\ell = \infty$ and $\sum_{\ell=0}^\infty \lambda_\ell^2 < \infty$, with $\lambda_0 = 1$ for initialization. This is essentially a Monte Carlo version of the generic online EM algorithm of Cappé and Moulines (2009); the main difference is that the integrals are w.r.t. to the target distribution instead of w.r.t. the empirical distribution of some observations as in classical EM.

As showed by Douc et al. (2008), it holds that

$$\frac{1}{N} \sum_{i=1}^N \tilde{\omega}_i \xrightarrow{\mathbb{P}} \iint \nu(\mathbf{x}) l(\mathbf{x}, \tilde{\mathbf{x}}) d\mathbf{x} d\tilde{\mathbf{x}}$$

as N tends to infinity; therefore, relying on the weighted sample $\{(I_j^{[\ell]}, \tilde{\mathbf{X}}_j^{[\ell]})\}_{j=1}^{N^\ell}$ produced at the current iteration, we use

$$\mathcal{C}^\ell := \begin{cases} \frac{1}{N^0} \sum_{i=1}^{N^0} \tilde{\omega}_i^{[0]}, & \text{for } \ell = 0, \\ (1 - \lambda_\ell) \mathcal{C}^{\ell-1} + \lambda_\ell \frac{1}{N^\ell} \sum_{i=1}^{N^\ell} \tilde{\omega}_i^{[\ell]}, & \text{for } \ell \geq 1, \end{cases}$$

as normalizing constant of the importance sampling estimate. We then estimate, for any function F , the expectation $\mathbb{E}_{\mu^{\text{aux}}}[F(\boldsymbol{\theta}, I, \tilde{\mathbf{X}})]$ by

$$\hat{\mathbb{E}}_{\mu^{\text{aux}}}^{[\ell]} [F(\boldsymbol{\theta}, I, \tilde{\mathbf{X}})] := (1 - \lambda_\ell) \hat{\mathbb{E}}_{\mu^{\text{aux}}}^{[\ell-1]} [F(\boldsymbol{\theta}^{\ell-1}, I, \tilde{\mathbf{X}})] + \lambda_\ell \sum_{i=1}^{N^\ell} \frac{\tilde{\omega}_i^{[\ell]}}{N^\ell \mathcal{C}^\ell} F(\boldsymbol{\theta}; I_i^{[\ell]}, \tilde{\mathbf{X}}_i^{[\ell]}). \quad (4.1)$$

Obviously this stochastic approximation of the normalizing constant does not guarantee that the renormalized weights sum to 1, and hence leads to approximation of a probability measure with a finite measure. However, this does not not impede the convergence of the stochastic approximation. Moreover, the weights of the final sample obtained from the adapted kernel are renormalized to sum to one: the final approximation therefore remains a probability distribution.

Past the initial step $\ell = 0$, which provides the first optimized fit, the numbers N^ℓ of particles sampled at each iteration $\ell \geq 1$ can be noticeably small. A

convergence proof of this stochastic approximation algorithm will be the topic of the aforementioned companion paper and is rooted in the analysis of generic Online EM (Cappé and Moulines, 2009), *stochastic approximation EM* (SAEM, Delyon et al., 1999; Kuhn and Lavielle, 2004; Andrieu et al., 2005) and *Monte Carlo EM* (MCEM, Fort and Moulines, 2003). In the present paper we limit ourselves to numerical illustrations of the proposed scheme; see the next section.

Remark 4.1 (Proposal kernel at the initial iteration). Finally, we may choose to draw the particles of the first step from a distribution different from π_{θ^0} in order to avoid relying on an arbitrary initial choice which is often quite poor. When operating on state-space models, as in the examples presented in the next section, we obtain the first sample $\{(I_i^{[0]}, \tilde{\mathbf{X}}_i^{[0]})\}_{i=1}^{N^0}$ using the prior kernel—the simplest default choice.

5 Applications to nonlinear state-space models

SMC methods can be successfully applied to optimal filtering in *state-space models*. A state-space model is a bivariate process $(X_k, Y_k)_{k \geq 0}$, where $X = (X_k)_{k \geq 0}$ is an unobserved Markov chain on some state-space $\mathbf{X} \subseteq \mathbb{R}^d$ and $(Y_k)_{k \geq 0}$ is an observation process taking values in some space $\mathbf{Y} \subseteq \mathbb{R}^p$. We denote by Q and π_0 the transition density and initial distribution of X , respectively. Conditionally on X , the observations are assumed to be conditionally independent with the conditional distribution G of a particular Y_k depending on the corresponding X_k only. We will assume that Q and G have densities q and g , respectively, with respect to Lebesgue measure, i.e.

$$\mathbb{P}(X_{k+1} \in A | X_k) = Q(X_k, A) = \int_A q(X_k, x) \, dx \quad (5.1)$$

and

$$\mathbb{P}(Y_k \in B | X_k) = G(X_k, B) = \int_B g(X_k, y) \, dy. \quad (5.2)$$

For simplicity, we have assumed that the Markov chain X is time-homogeneous and that the distribution G does not depend on k ; however, all developments that follow may be extended straightforwardly to the time-inhomogeneous case. The optimal filtering problem consists of computing, given a fixed record $Y_{0:n} := (Y_0, \dots, Y_n)$ of observations, the *filter distributions* $\phi_k(A) := \mathbb{P}(X_k \in A | Y_{0:k})$ for $k = 0, 1, \dots, n$. As the filtering distributions satisfy the nonlinear recursion

$$\phi_k(x_{k+1}) = \frac{\int g(x_{k+1}, y_{k+1}) q(x_k, x_{k+1}) \phi_k(x_k) \, dx_k}{\iint g(x_{k+1}, y_{k+1}) q(x_k, x_{k+1}) \phi_k(x_k) \, dx_k \, dx_{k+1}},$$

the optimal filtering problem can be perfectly cast into the sequential importance sampling framework (1.1) with $\Xi = \tilde{\Xi} \equiv \mathbf{X}$, $\nu \equiv \phi_k$, $\mu \equiv \phi_{k+1}$, and

$l(\mathbf{x}, \tilde{\mathbf{x}}) \equiv g(\tilde{\mathbf{x}}, y_k)q(\mathbf{x}, \tilde{\mathbf{x}})$. Consequently, a particle sample $\{(\mathbf{X}_i, \omega_i)\}_{i=1}^N$ approximating ϕ_k can be transformed into a sample $(\tilde{\mathbf{X}}_i, \tilde{\omega}_i)_{i=1}^N$ approximating the filter ϕ_{k+1} at the subsequent time-step by (1) drawing indices $\{I_i\}_{i=1}^N$ multinomially with weights proportional to $(a_i \omega_i)_{i=1}^N$, (2) simulating particle positions $\tilde{\mathbf{X}}_i \sim r(\mathbf{X}_{I_i}, \cdot)$, and (3) assigning each updated particle $\tilde{\mathbf{X}}_i$ the importance weight $\tilde{\omega}_i = g(\tilde{\mathbf{X}}_i, y_{k+1})q(\mathbf{X}_{I_i}, \tilde{\mathbf{X}}_i)/(a_{I_i}r(\mathbf{X}_{I_i}, \tilde{\mathbf{X}}_i))$. Here $(a_i)_{i=1}^N$ is a set of nonnegative adjustment multipliers and the proposal r is a Markov transition density. In the following examples we illustrate how the proposal r can be designed adaptively using mixtures of experts.

5.1 Multivariate linear Gaussian model

We start by considering a simple multivariate linear Gaussian model. In this toy example, the optimal kernel L^* and the optimal adjustment weight function a^* are available on closed form. This makes it possible to compare our algorithm to an exact reference. The optimal kernel does not belong to our family of mixture of experts. In the model under consideration:

- The distribution ν on $\mathbf{X} = \mathbb{R}^2$ is bimodal with density $\nu(\mathbf{x}) = 0.5\mathcal{N}_2(\mathbf{x}; (0, 1)^T, \Sigma) + 0.5\mathcal{N}_2(\mathbf{x}; (0, -1)^T, \Sigma)$, where $\Sigma = 0.1I_2$. The two modes are hence well separated.
- The prior kernel density $q(\mathbf{x}, \tilde{\mathbf{x}}) = 0.5\mathcal{N}_2(\tilde{\mathbf{x}}; \Lambda_1 \bar{\mathbf{x}}, \Sigma) + 0.5\mathcal{N}_2(\tilde{\mathbf{x}}; \Lambda_2 \bar{\mathbf{x}}, \Sigma)$ is a mixture of multivariate Gaussian distributions with regression matrices


$$\Lambda_1 = \begin{pmatrix} 1 & 0 & 1 \\ 0 & 1 & 1 \end{pmatrix}, \quad \Lambda_2 = \begin{pmatrix} 1 & 0 & 1 \\ 0 & 1 & -1 \end{pmatrix}.$$

- Each Y_k , taking values in $\mathbf{Y} = \mathbb{R}^2$, is a noisy observation of the corresponding hidden state X_k with local likelihood $g(\tilde{\mathbf{x}}, y) = \mathcal{N}_2(y; \tilde{\mathbf{x}}, \Sigma_Y)$, where $\Sigma_Y = 0.1I_2$.

In this first example, we consider one single step of the particle filter, with initial particles $\{\mathbf{X}_i\}_{i=1}^N$ sampled exactly from the initial distribution. The observation Y_k is chosen to be $Y_k = (1, 0)^T = \Lambda_2(0, 1)^T = \Lambda_1(0, -1)^T$. The ancestor of the hidden state around which this observation is centered is equally likely to belong to any of the two components of the original distribution, depending on which (unknown) component of the prior kernel was used to sample the new state. Even though the optimal kernel L^* does not belong to our family of experts, it is still a mixture whose weights are highly dependent of location of the ancestor, and we can expect our algorithm to adapt accordingly.

The histogram in Figure 1a of the importance weights of the particle swarm produced using the prior kernel shows that 50% of these weights are nearly equal to zero, and the remaining ones are spread over a large range of values. This

is confirmed by looking at the corresponding curve of proportions in Figure 1b, showing that 80% of the total mass is carried by only 25% of the particles and 99% of the total mass by 40% of the particles.



(a) Histogram of the importance weights obtained for prior kernel and the final adapted kernel in the linear Gaussian model. Weights are re-normalized and multiplied by the size of the sample. In exact sampling, all weights would be equal to unity.

(b) *Curve of proportions*: proportion of the particles, sorted by decreasing order of importance weight, against proportion of the total mass, for the prior kernel and for several iterations of the adaptation algorithm.

Figure 1: Evolution of the distribution of the importance weights before and after adaptation.

One single iteration of our SA-based adaptation scheme described in Section 4 is sufficient to improve those proportions to 80% of the mass for 40% of the particles and 99% of the mass for 55% of the particles; see the corresponding curve of proportion in Figure 1b. After 10 such iterations, the corresponding curve of proportions in Figure 1b shows that close to maximum improvement has been reached: the first few steps of the optimization are enough to bring significant improvement. The histogram in Figure 1a of the weights obtained at the final iteration shows how the weights are concentrated around unity, the value that corresponds to sampling with the optimal kernel.

We display in Figure 2a the KLD (1.6) between the fit and the target (estimated by means of a Monte Carlo approximation with a large sample size), and the KLD for the proposal with prior kernel and for the proposal with optimal kernel with uniform adjustment weights — the optimal adjustment weights are near uniform in this example. As mentioned earlier, the KLD decreases very fast, most of the improvement being obtained in the first few iterations. Figure 2b compares the evolution of the KLD for several step sizes covering four order of magnitude. As the step size is increased, the algorithm converges faster, although less smoothly: this has no practical impact as we are only looking for a good

proposal kernel in an importance sampling setting, not for the exact optimal one.

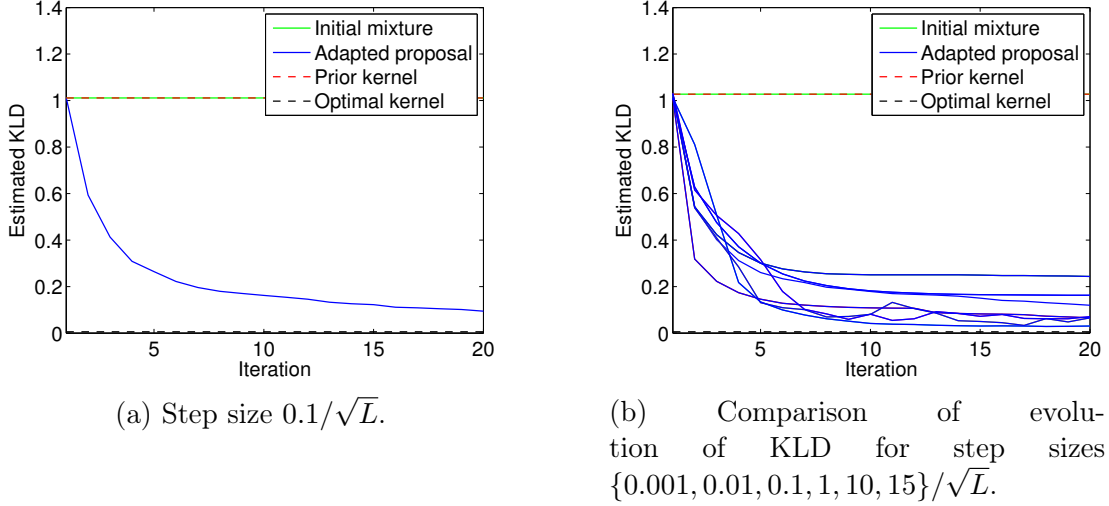


Figure 2: Evolution of the KLD over $L = 20$ iterations of adaptation for the LG model.

5.2 Bessel process observed in noise

As a more intricate example, we consider here optimal filtering of the Brownian motion underlying a Bessel process observed in noise, also known as range-only filtering of a Gaussian random walk. We let the state process evolve in the plane for visualization purposes. More specifically, the state-space model is given by

$$\begin{aligned} X_{k+1} &= X_k + V_k, \\ Y_k &= \|X_k\| + W_k, \end{aligned}$$

where $V_k \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}_2(0, \Sigma_x)$ and $W_k \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(0, \sigma_y^2)$ are mutually independent, and $\|X\| = \sqrt{\sum_{i=1}^2 X_{(i)}^2}$ is the L_2 norm on the state space $\mathbf{X} := \mathbb{R}^2$. With the notations of (5.1) and (5.2), we define by $q(\mathbf{x}, \tilde{\mathbf{x}}) := \mathcal{N}_2(\tilde{\mathbf{x}}; \mathbf{x}, \Sigma_x)$ the density of the prior kernel and by $g(\tilde{\mathbf{x}}, y) := \mathcal{N}(y; \|\tilde{\mathbf{x}}\|, \sigma_y^2)$ the local likelihood. We ensure a diffuse prior kernel and informative observations by setting $\Sigma_x = I_2$ and variance $\sigma_y^2 = 0.01$. As the hidden state is observed range-only, the state equation provides most of the information about the bearing while the local likelihood is invariant under rotation of the state around the origin. This induces a variety of nonlinear shapes of the optimal kernel depending on the location of the ancestor—see the three top rows of Figure 3.

The original weighted sample $\{(\mathbf{X}_i, \omega_i)\}_{i=1}^N$ consists of a sample of $N = 20,000$ i.i.d. particles drawn from the prior distribution $\nu = \mathcal{N}_2((0.7, 0.7), 0.5I_2)$ of the

state at time 0 (and hence having uniform weights). The first observation is $Y_1 = 1.0$. We initialize the algorithm with $N_1 = 1,000$ particles $\{(\tilde{\mathbf{X}}_i, \tilde{\omega}_i)\}_{i=1}^{N_1}$ using the prior kernel. The resulting cloud, along with the 100 particles with the largest weights and their ancestors, is plotted in Figure 4 (top row). The support of the proposal distribution is over-spread: most of the particles have negligible weights, and a few particles have comparatively large weights: 80% of particles have null weight. This is confirmed by the curve of proportions in Figure 5 (left): only 20% of the proposed particles carry the total mass. Adaptation of the proposal kernel is thus relevant. Here again, adaptation of the adjustment weights is not required, as the ancestors of the particles with highest importance weights are not located in any specific region of the space.

Based on these 1,000 particles, the first iteration of the adaptation is done using conditional probabilities from the initial fit displayed in Figure 3 (fourth row) whose components are chosen to be independent of the ancestor, i.e. only the constant term is non-null. The resulting kernel is plotted in Figure 3 (fifth row). We use it to propose 200 new particles, serving as importance sampling approximation of μ^{aux} at iteration 2. After 30 such iterations, the adapted kernel visible in Figure 3 (last row) is visually very close to the optimal kernel. Note the impact of the location of the ancestor on the un-normalized transition kernel, whose mass shifts on a circle, and how the adapted kernel shifts accordingly. Figure 4 (bottom row) shows that adaptation concentrates the particles on the support of the target distribution, thus narrowing the span of the importance weights and balancing the contributions of all particles.

Most importantly, a very small number of iterations suffices to achieve significantly more uniformly distributed importance weights, i.e. significantly lower KLD: Figure 5 (right) shows how the KLD drops after the first 2 iterations and then stabilizes near null, while the curve of proportions in Figure 5 (left) shows that the distribution of the weights is essentially unchanged past the first few iterations.

As a final look into the convergence, Figure 6 displays the evolution of all the estimated parameters over the 30 iterations, confirming that the fit stabilizes after a few steps: the result of the very first couple of iterations could serve as a more efficient proposal than the prior kernel. The parameters of the logistic weights $\underline{\beta}^\ell$ are the slowest to stabilize due to the stochastic gradient used to palliate for the lack of closed-form update as mentioned in Section 3.

We now iterate the same procedure over 50 time-steps and compare the adaptive filter with the plain bootstrap filter. The forgetting of the initial distribution of the random walk entails that the filter distribution flow converges to a distribution symmetric under rotation and supported simply on a ring centered at $(0, 0)$. To connect with widespread measures of efficiency discussed in Cornebise et al. (2008), we compare in Figure 7 the negated Shannon entropy (lower is better) of the importance weights and relative effective sample size (in percentage of the total sample; higher is better) of the two filters. The negated Shannon entropy




Figure 3: Visual evolution of the target and adapted kernels in the Bessel model, for 3 distinct ancestors in different regions, for observation $y = 1.0$.

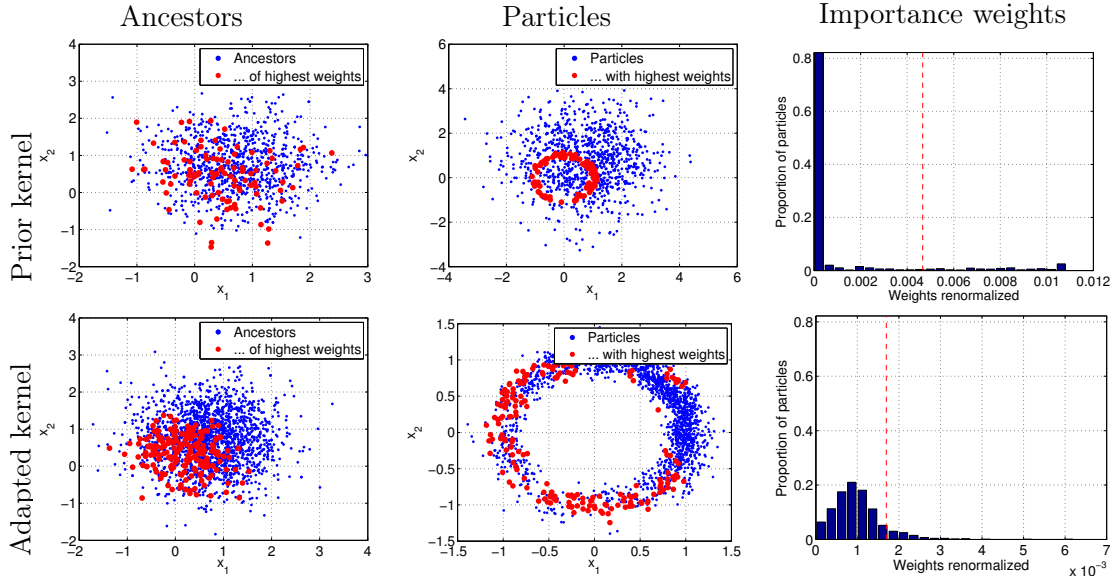


Figure 4: Comparison of 1,000 and 200 particles from the prior and final adapted kernels, respectively, with the 10% of largest weights plotted in red, along with their ancestors and the distribution of their importance weights.




Figure 5: Improvement of the adaption over 30 iterations. The proportions plot (left) and the KLD (right) dramatically improve over the very first iterations. While the prior kernel puts 90% of the mass on 15% of the particles, adaptation increases it to 70% of the particles in one step and stabilizes to 80% after a very few iterations.

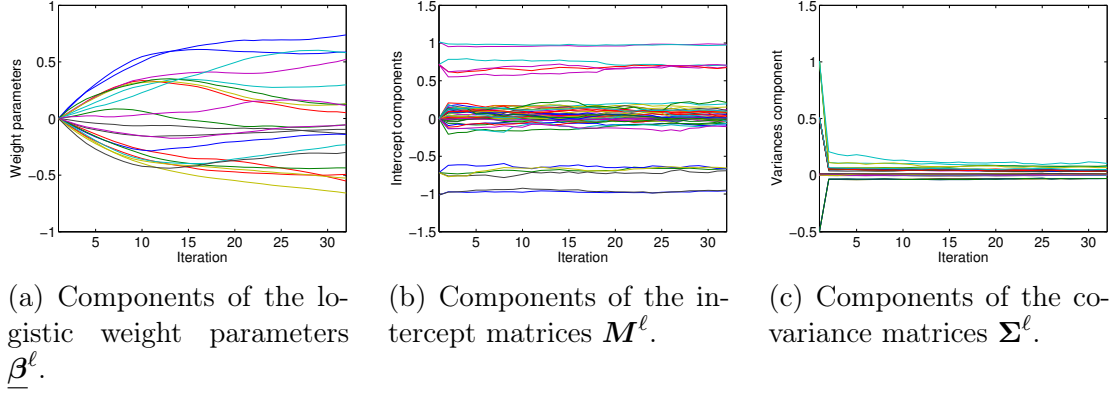


Figure 6: Parameters θ^ℓ of the adapted kernel over 30 iterations of the algorithm: practically, convergence is achieved after a few steps only.

shows improvement by our adaptive filter, which is not surprising in the light of the convergence results of (Cornebise et al., 2008, Theorem 1) recalled in Section 1: it converges to the same KLD between the same distributions when the number of particles tends to infinity. More remarkably, in this example, the relative effective sample size also shows an improvement, while the same convergence results prove it to be linked to the chi-square divergence instead of the KLD. While a first intuition would make sense of improving the CSD by minimizing the KLD, this is not systematically the case.

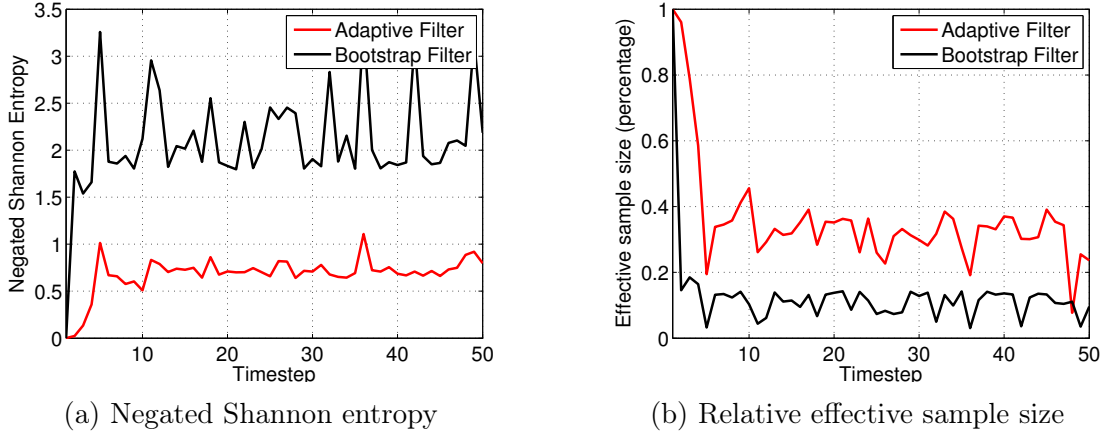


Figure 7: Evolution of the effective sample size and the entropy between target and proposal for the bootstrap filter and the adaptive filter. As intended, the adaptive filter leads to higher effective sample size and lower negated Shannon entropy.

5.3 Multivariate tobit model

We now briefly illustrate how far adaptation of the proposal kernel goes—and where it stops. Consider a partially observed multivariate dynamic *tobit* model

$$\begin{aligned} X_k &= A X_{k-1} + U_k , \\ Y_k &= \max(B^T X_k + V_k, 0) , \end{aligned}$$

where $\mathbf{X} = \mathbb{R}^2$, A is a 2×2 invertible matrix, and $B = (1, 1) \in \mathbb{R}^2$, so that Y_k takes values in \mathbb{R} . The random variables U_k and V_k are Gaussian white noises with variances Σ_U and σ_v^2 , respectively. The observation process consists of noisy observations of the sum of its components of the hidden states. In addition, the observations are left-censored. We have an original swarm of $N = 20,000$ particles distributed according to $X_0 \sim \mathcal{N}_2((1, 1)^T, 10I_2)$ and we set $Y_1 = 0$. The local likelihood is hence null above the line $\Delta = \{x \in \mathbb{R}^2 : B^T x = 0\}$, and constant below, with a narrow transition region, as displayed in Figure 8 (second row). The prior kernel displayed in Figure 8 (first row) can have most of its mass out of the high-likelihood regions, depending on the ancestor. Figure 8 (third row) illustrates the un-normalized optimal transition kernel for three reference ancestors, showing how the match between the supports of the proposal kernel and of the local likelihood varies depending on the position of the ancestor particle relatively to the line $A^{-1}\Delta$. Hence half of the original particles have very small adjustment multiplier weights.

Adapting the proposal kernel will hence require at least two components, and will not lead to perfect adaptation, as can be seen on the fit after 10 iterations displayed in Figure 8 (last row). The ancestor $(1, 1)$ is the center of the ancestor sample, un-normalized optimal kernel is the prior kernel truncated in its middle; $(-3, -1)$ is the bottom left of the ancestors sample, the un-normalized optimal kernel almost matches the prior, save for a truncation in the upper right tail; $(9, 5)$ is the top right of the ancestor sample, the un-normalized optimal kernel differs widely from the prior kernel, as only the very far tails of the latter match non-null local likelihood.

Finally, without getting into details, we make our point through Figure 9 showing again how the KLD drops in the first iteration for families of both Gaussian distributions and Student's t -distributions. It then stabilizes close to the minimum achieved by the optimal kernel with uniform adjustment weights, which is not negligible. The Student's t -distribution allows heavier tails, at the price of a higher attainable lower bound on the KLD.

6 Future work and Conclusion

Relying on the results of [Cornebise et al. \(2008\)](#), we have built new algorithms approximating the so-called optimal proposal kernel at a given time step of an




Figure 8: Visual evolution of the adapted kernel, for 3 ancestors in different regions for the Tobit model.

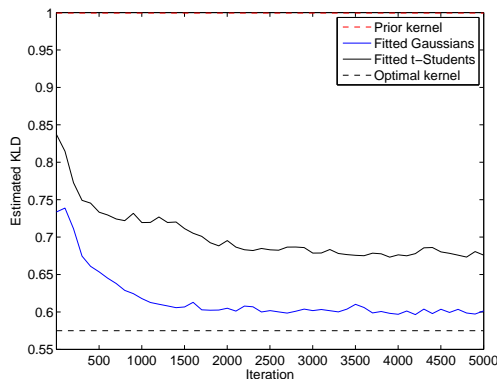


Figure 9: Comparison of the evolution of the KLD for the Gaussian experts and the Student’s t -distributed expert over 5,000 iterations of the algorithm.

auxiliary particle filter by means of minimization of the KLD between the auxiliary target and instrumental distributions of the particle filter. More specifically, the algorithm fits a weighted mixture of integrated curved exponential distributions with logistic weights to the auxiliary target distribution by minimizing the KLD between the two using a Monte Carlo version of the online EM method proposed by [Cappé and Moulines \(2009\)](#).

In addition, we have applied successfully this relatively simple algorithm to optimal filtering in state-space models; indeed, running the stochastic approximation-based adaptation procedure for only a few iterations leveled off significantly the distribution of the total weight mass among the particles also for models exhibiting very strong nonlinearity. Thus, adding the adaptation step to an existing particle filter implementation implies only limited computational demands.

A proof of convergence, along the lines mentioned in Section 4, of the algorithm is currently in progress. In addition, we investigate at present the possibility of extending the approach to comprise adaptation also of the adjustment multipliers.

References

- C. Andrieu and É. Moulines. On the ergodicity properties of some adaptive MCMC algorithms. *The Annals of Applied Probability*, 16:1462–1505, 2006.
- C. Andrieu, É. Moulines, and P. Priouret. Stability of stochastic approximation under verifiable conditions. *SIAM Journal on Control and Optimization*, 44(1): 283–312, 2005.
- C. Andrieu, A. Doucet, and R. Holenstein. Particle markov chain monte carlo

- methods. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 72(3):269–342, 2010.
- O. Cappé and É. Moulines. On-line expectation–maximization algorithm for latent data models. *J. Roy. Statist. Soc. Ser. B*, 71(3):593–613, 2009.
- O. Cappé, E. Moulines, and T. Rydén. *Inference in Hidden Markov Models*. Springer, 2005.
- O. Cappé, R. Douc, A. Guillin, J. M. Marin, and C. P. Robert. Adaptive importance sampling in general mixture classes. *Statistics and Computing*, 18(4):447–459, 2008.
- K. Chen, L. Xu, and H. Chi. Improved learning algorithms for mixture of experts in multiclass classification. *Neural Networks*, 12(9):1229–1252, 1999.
- J. Cornebise. *Adaptive Sequential Monte Carlo Methods*. PhD thesis, Université Pierre et Marie Curie – Paris 6, June 2009.
- J. Cornebise, E. Moulines, and J. Olsson. Adaptive methods for sequential importance sampling with application to state space models. *Statistics and Computing*, 18(4):461–480, 2008.
- P. Del Moral, Arnaud Doucet, and Ajay Jasra. Sequential monte carlo samplers. *J. Roy. Statist. Soc. Ser. B*, 68(3):411, 2006.
- B. Delyon, M. Lavielle, and E. Moulines. On a stochastic approximation version of the EM algorithm. *Ann. Statist.*, 27(1), 1999.
- R. Douc, É. Moulines, and J. Olsson. Optimality of the auxiliary particle filter. *Probability and Mathematical Statistics*, 28(2), 2008. To appear.
- A. Doucet, S. Godsill, and C. Andrieu. On sequential Monte Carlo sampling methods for Bayesian filtering. *Statistics and Computing*, 10(3):197–208, 2000.
- A. Doucet, N. De Freitas, and N. Gordon, editors. *Sequential Monte Carlo Methods in Practice*. Springer, New York, 2001.
- G. Fort and E. Moulines. Convergence of the Monte Carlo expectation maximization for curved exponential families. *Ann. Statist.*, 31(4):1220–1259, 2003.
- N. Gordon, D. Salmond, and A. F. Smith. Novel approach to nonlinear/non-Gaussian Bayesian state estimation. *IEE Proc. F, Radar Signal Process.*, 140:107–113, 1993.
- H. Haario, E. Saksman, and J. Tamminen. An adaptive Metropolis algorithm. *Bernoulli*, 7(2):223–242, 2001.

- M. Jordan and R. Jacobs. Hierarchical mixtures of experts and the EM algorithm. *Neural computation*, 6:181–214, 1994.
- M. Jordan and L. Xu. Convergence results for the em approach to mixtures of experts architectures. *Neural Networks*, 8(9):1409–1431, 1995. ISSN 0893-6080.
- E. Kuhn and M. Lavielle. Coupling a stochastic approximation version of EM with an MCMC procedure. *ESAIM Probab. Statist.*, 8:115–131, 2004.
- C. Liu and D.B. Rubin. ML estimation of the t distribution using EM and its extensions, ECM and ECME. *Statistica Sinica*, 5(1):19–39, 1995.
- J.S. Liu. *Monte Carlo Strategies in Scientific Computing*. Springer, New York, 2001.
- G.J. McLachlan and T. Krishnan. *The EM algorithm and extensions, second edition*. Wiley New York, 2008.
- M.-S. Oh and J. O. Berger. Integration of multimodal functions by Monte Carlo importance sampling. *J. Amer. Statist. Assoc.*, 88(422):450–456, 1993. ISSN 0162-1459.
- D. Peel and G.J. McLachlan. Robust mixture modelling using the t distribution. *Statistics and Computing*, 10(4):339–348, 2000.
- M. K. Pitt and N. Shephard. Filtering via simulation: Auxiliary particle filters. *J. Am. Statist. Assoc.*, 94(446):590–599, 1999.
- R.E. Quandt and J.B. Ramsey. A new approach to estimating switching regressions. *J. Am. Statist. Assoc.*, 67:306–310, 1972.
- B. Ristic, M. Arulampalam, and A. Gordon. *Beyond Kalman Filters: Particle Filters for Target Tracking*. Artech House, 2004.
- G. O. Roberts and J. S. Rosenthal. Coupling and ergodicity of adaptive markov chain monte carlo algorithms. *Journal of Applied Probability*, 44:458–475, 2007.
- G. O. Roberts and J. S. Rosenthal. Examples of adaptive MCMC. *Journal of Computational and Graphical Statistics*, 18(2):349–367, 2009.
- R. Y. Rubinstein and D. P. Kroese. *The Cross-Entropy Method*. Springer, 2004.
- R. Van der Merwe and E. Wan. Sigma-point kalman filters for probabilistic inference in dynamic state-space models. In *Proceedings of the Workshop on Advances in Machine Learning*, Montreal, Canada, June 2003.
- R. Van der Merwe, A. Doucet, N. De Freitas, and E. Wan. The unscented particle filter. In T. K. Leen, T. G. Dietterich, and V. Tresp, editors, *Adv. Neural Inf. Process. Syst.*, volume 13. MIT Press, 2000.